4.6 A 20Gb/s Forwarded Clock Transceiver in 90nm CMOS

B. Casper[1], J. Jaussi[1], F. O'Mahony[1], M. Mansuri[1], K. Canagasaby[2], J. Kennedy[1], E. Yeung[2], R. Mooney[1]

[1]Intel Labs, Hillsboro, OR
[2]Intel, Santa Clara, CA

Future microprocessor platforms will require system-level optimization of the I/O to minimize cost and maximize aggregate bandwidth. Critical parameters such as silicon area, power, testability, and off-chip interconnect quality must be properly balanced to maximize the I/O performance versus cost ratio. For example, there is a fundamental tradeoff between clock quality and equalizer effectiveness [1]. Producing precision RX and TX clocks and a sensitive RX may impact the power and area of some circuits, but it could allow the use of simple, low-power linear equalizers to minimize the overall link power. Additionally, these equalizers will become much more effective by limiting near-end crosstalk and stubbed backplane (BP) via length. To demonstrate this system-level optimization effort, we have developed a 20Gb/s forwarded clock I/O system intended for a wide parallel link with small area and low power.

The TX architecture shown in Fig. 4.6.1 uses a 4-tap, 6-bit coefficient resolution equalizer to compensate for channel losses up to -20dB at 10GHz. The clock is generated by a PLL with an LC-VCO and passively distributed to provide very low clock-to-data TX jitter which is required for effective forwarded clock systems. While long-term phase noise has little consequence for the clocking method we have developed, the interaction between high-frequency jitter and ISI limits performance and has been alleviated with the use of a 500kHz loop BW PLL to limit reference-clock jitter transfer. Additionally, we used supply-insensitive clock buffers and a VCO with local duty-cycle correction (DCC) to ease the internally generated high-frequency jitter. The TX equalizer (TXEQ) is based on a 4-way interleaved lookup table [2] with each section operating at 5GHz. Since the output of the 16-bit by 6-bit lookup table is static, we minimized area by sharing the coefficient signals across all of the 16 to 1 MUXES. The MUX select lines are driven by a 4-bit sequence of the TX pattern. A lookup-table-based equalization provides the flexibility to place the cursor tap at any one of 4 positions and allows both linear and non-linear forms of equalization. To ensure the 6-bit TX DAC architecture has sufficient linearity and delay accuracy, DAC segments were carefully matched. Rather than using 64 individual segments for the 2 to 1 serializer, predriver and output stage, the primitive segment is chosen to include 1 serializer, 1 predriver, and 4 output stages. This architecture minimizes the number of DAC segments to 17, which corresponds to a reduction in power and area while producing an output INL of better than 0.12LSB.

A 5GHz forwarded clock is sent to an RX-side on-die transmission-line clock distribution that is impedance matched to the off-chip interconnect. This alleviates the need for RX global clock buffering and the associated power and jitter penalty. Figure 4.6.2 demonstrates the locally received clock being sent through AC coupling, buffers and DCC before being transmitted to the static alignment mechanism. The AC coupling rejects sub-GHz common-mode noise and translates the clock levels for compatibility with the NMOS input symmetric-load delay cells used in the buffers and VCDLs. The 8-stage DLL and interpolator enable deskew of each data lane to optimally sample the eye. The RX sampling phases of 0˚, 90˚, 180˚, and 270˚ are generated using a VCDL that is matched and delay-slaved to the DLL. The data is received with an RX continuous-time LE (RXEQ) that is based on adjustable source degeneration of a differential pair. The RXEQ buffers the data to a 4-way interleaved comparator set, each operating at 5Gb/s. Offset and delay-mismatch compensation are achieved through the static calibration of the 4 comparators. To enable diagnostic features such as waveform capture, BER eye diagram plotting and circuit uncertainty characterization [3], wide range offset margining is provided through a differential 8-bit DAC that sums into the output nodes of the RXEQ.

The I/O link was characterized over both FR4 single-board (2 to 7 inches) and BP (7 to 34 inches) channels. The BP channel consisted of two 2-inch FR4 traces and connectors linked by a Rogers or FR4 BP. Figure 4.6.3 demonstrates the achievable data rates as a function of interconnect topology and length. All channels included an LGA package with Socket-Ts at both the TX and RX. The on-die measured (ODM) eye diagram is shown for the 7-inch link at 20Gb/s in Fig. 4.6.3.

We used on-die waveform capture to demonstrate the 20Gb/s single-bit pulse responses in Fig. 4.6.4 (a). Each equalizer setting was optimized for the respective maximum data rate (shown in Fig. 4.6.6) and not necessarily 20Gb/s. Even though the TXEQ pulse response appears to have a higher signal to ISI ratio than the combined equalizer response, it actually provides 4% lower link performance. Combining equalization methods exploits the inherent advantages of TXEQ and RXEQ. While TXEQ allows high-order compensation of channel losses and reflections, high-frequency jitter at the TX can severely degrade equalizer performance [1]. Conversely, RXEQ does not amplify high-frequency TX jitter, but there is generally less flexibility regarding the cancellation of ISI. The RXEQ tuning is represented through low-frequency (DC) gain and the source degeneration pole's high-frequency (AC) gain, as shown in the example of Fig. 4.6.4 (b). Figure 4.6.4 (c) shows the TXEQ normalized coefficient sensitivity to maximum data rate with a fixed setting of RXEQ. Figure 4.6.4 (d) demonstrates the RXEQ sensitivities with constant TXEQ coefficients.

As previously described, jitter performance is one of the most sensitive parameters to link rates. DCC at the TX plays a key role in achieving 20Gb/s, as shown in Fig. 4.6.5 (a). To show the effectiveness of the forwarded clock system, we tested the tolerance of the link to deterministic jitter on the TX clock at 13.5Gb/s. As shown in Fig. 4.6.5 (b), the measured jitter tolerance we achieved at 40MHz exceeds the OC-192 specification by more than 25 times. The most important measure of jitter in a forwarded clock I/O is the link timing uncertainty using ODM methods. The distribution in Fig. 4.6.5 (c) is generated by sweeping the calibrated interpolator across an ISI-free data edge while sampling a comparator. Other applicable jitter measurements using external test equipment are shown in the jitter summary table in Fig. 4.6.5. Link performance is summarized in Fig. 4.6.6 and the die microphotograph is shown in Fig. 4.6.7.

Acknowledgements:
We thank T. Hollis, M. Maeng, S. Venkataraman, S. Jain, P. Iyer, R. Rajendran, M. Parthasarthy, J. Howard, D. Finan, H. Wilson, S. Dighe, L. Rathnam, H. Nguyen, K. Truong, C. Parsons, for prototype assistance, V. Prokofiev, P. Parmar, C. Roberts, D. Casper for interconnect and test support, C. Sweeney and Intel Fab 24 for silicon fabrication.

References:
[1] V. Stojanoviç and M. Horowitz, "Modeling and Analysis of High Speed Links," *IEEE CICC Proc.*, pp. 589–594, 2003.
[2] W. Dally and J. Poulton, "Transmitter Equalization for 4Gb/s Signaling," *presented at the Hot Interconnects Conf.*, 1996.
[3] B. Casper et al., "8 Gb/s SBD Link with On-Die Waveform Capture," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2111–2120, Dec., 2003.

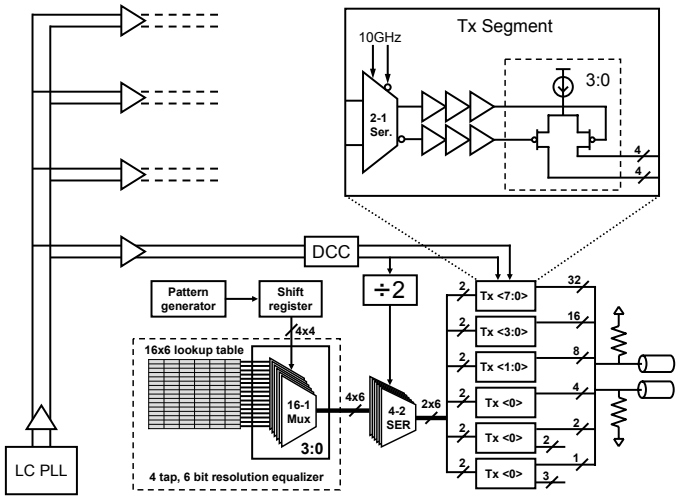

Figure 4.6.1: TX architecture and circuits.

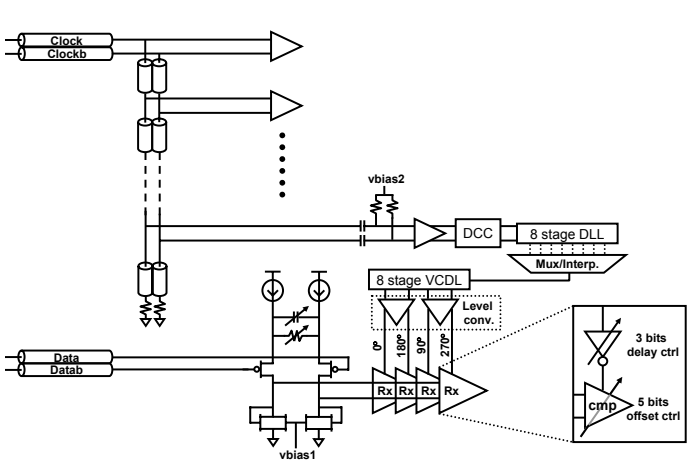

Figure 4.6.2: Forwarded clock and RX architecture.

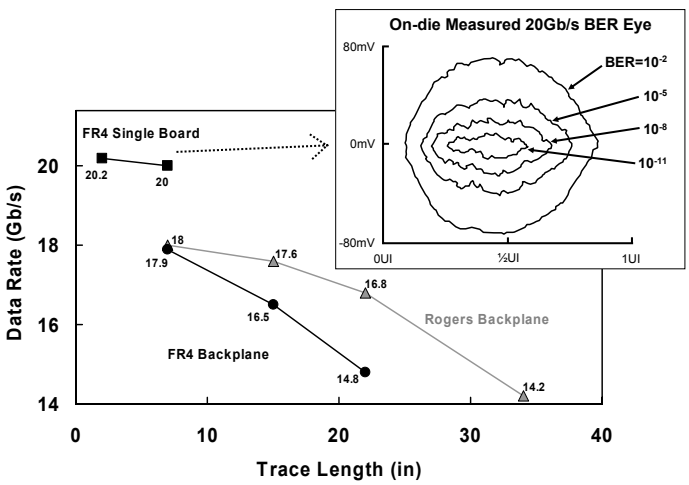

Figure 4.6.3: Data rate vs. channel length at 10^{-12} BER using a 32b LFSR.

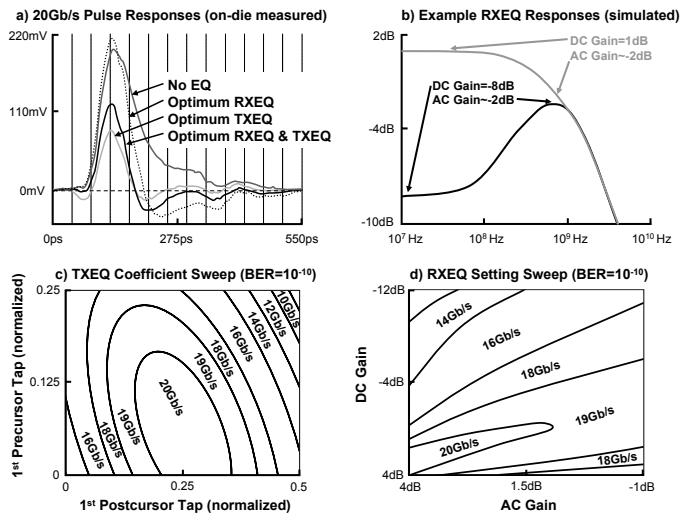

Figure 4.6.4: Measured TX and RX equalization performance for 7" FR4 link.

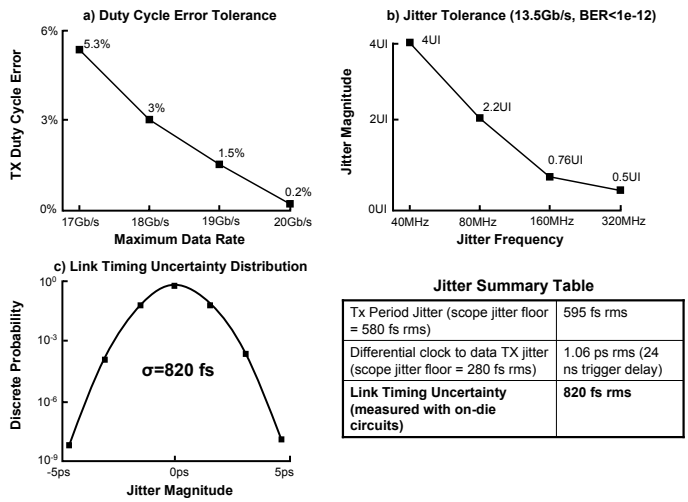

Jitter Summary Table

Tx Period Jitter (scope jitter floor = 580 fs rms)	595 fs rms
Differential clock to data TX jitter (scope jitter floor = 280 fs rms)	1.06 ps rms (24 ns trigger delay)
Link Timing Uncertainty (measured with on-die circuits)	**820 fs rms**

Figure 4.6.5: Measured jitter performance.

Process Technology	90nm, 7 metal CMOS
Supply Voltage	1.2V
RX input-referred noise	1 mV rms
Link Timing Uncertainty	820 fs rms
Pad Capacitance	RX < 100fF TX < 400fF
20Gb/s Transceiver Power Dissipation (global clocking power amortized across 16 bits)	237mW (11.8 mW/Gb/s)
Transceiver Area	RX = 0.07mm^2 TX = 0.09mm^2
Die Area	5470µm by 5470µm
Tx DAC Linearity (6 bit resolution)	INL=0.12 LSB DNL=0.08 LSB
Maximum Data Rates vs. Equalization (7" FR4, 2 sockets and packages, -16dB loss at 10GHz, 32-bit LFSR, BER<10^{-12})	
Optimum TX and RX equalization	20.0Gb/s
Optimum TX equalization only	19.3Gb/s
Optimum RX equalization only	17.3Gb/s

Figure 4.6.6: Summary table.

Continued on Page 640

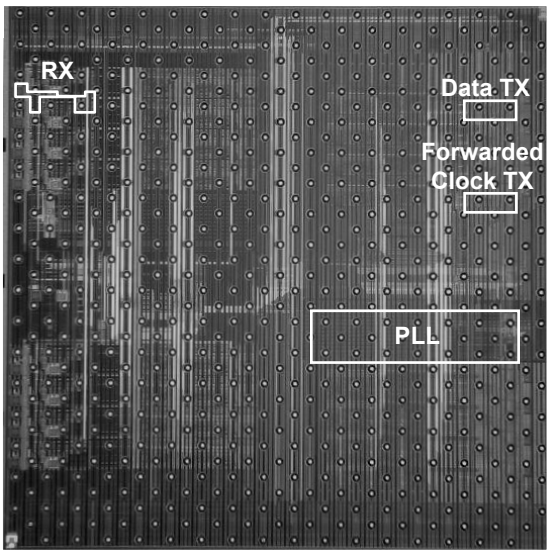

Figure 4.6.7: Die micrograph.

 1-4244-0079-1/06/$20.00